DRAFT - PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	ACS Enhanced Capability
DATE:	Introduced: 19 April 2018 Updated: 31 January 2019 Final Approval: TBD
AFFECTED DOCUMENT:	PCIe Specification
SPONSOR:	Intel Corporation Hewlett Packard Enterprise Fungible, Inc.

Part I

1. Summary of the Functional Changes

This ECN defines four new services under ACS for Downstream Ports, primarily to address issues when ACS redirect mechanisms are used to ensure that DMA Requests from Functions under the direct control of VMs are always routed correctly to the Translation Agent in the host. Three of the services provide redirect or blocking of Upstream Memory Requests that target areas not covered by other ACS services. The fourth service enables the blocking of Upstream I/O Requests, addressing a concern with VM-controlled Functions maliciously sending I/O Requests.

2. Benefits as a Result of the Changes

This change improves the level of isolation and protection provided by ACS, preventing accidental or malicious access of resources not covered by the current services in ACS. This especially improves the robustness of systems that use virtualization by ensuring that DMA Requests generated by VM-controlled Functions are always routed correctly to the Translation Agent and by removing potential attack vectors for Functions directly controlled by malicious VMs.

3. Assessment of the Impact

See 4 & 5 below.

4. Analysis of the Hardware Implications

This change only affects Switches and Root Ports that implement ACS Enhanced Capability features. All other hardware is unaffected.

5. Analysis of the Software Implications

Existing software is compatible with this ECN. No change is required. Updated software will receive the benefit of this ECN when enabling these new features.

6. Analysis of the C&I Test Implications

This change is compatible with existing tests.
Additional tests will be required to check the functionality added by this feature.

Part II

Detailed Description of the change

Make the following changes to Section 6.12:

…

ACS provides the following types of access control:

- ACS Source Validation ~~(V)~~
- ACS Translation Blocking ~~(B)~~
- ACS P2P Request Redirect ~~(R)~~
- ACS P2P Completion Redirect ~~(C)~~
- ACS Upstream Forwarding ~~(U)~~
- ACS P2P Egress Control ~~(E)~~
- ACS Direct Translated P2P ~~(T)~~
- ACS DSP Memory Target Access
- ACS USP Memory Target Access
- ACS Unclaimed Request Redirect
- ACS I/O Request Blocking

The specific requirements for each of these are discussed in the following section. ~~The letter in parenthesis following each type is the abbreviation for the associated capability and control bits defined in Section 7.7.7.~~

Make the following changes to Section 6.12.1.1:

6.12.1.1 ACS Downstream Ports

…

Completions are never affected by ACS Direct Translated P2P.

- ACS DSP Memory Target Access: must be implemented by Root Ports and Switch Downstream Ports that support ACS Enhanced Capability.

  ACS DSP Memory Target Access determines how an Upstream Request received by the Downstream Port's Ingress and targeting any Memory BAR Space[1] associated with an applicable Downstream Port is handled. The Request can be blocked, redirected, or allowed to proceed directly to its target. In a Switch, all Downstream Ports are applicable, including the one on which the Request was received. In a Root Complex, the set of applicable Root Ports is implementation specific, but always includes the one on which the Request was received.

[1] This includes any Memory Space allocated by an Expansion ROM Base Address register (BAR). This also includes any Memory Space allocated by EA entries with a BEI value of 0, 1, 7, or 8. See Section 7.8.5.2.

- ACS USP Memory Target Access: must be implemented by Switch Downstream Ports that support ACS Enhanced Capability; is not applicable to Root Ports.

  ACS USP Memory Target Access determines how an Upstream Request received by the Switch Downstream Port's Ingress and targeting any Memory BAR Space[2] associated with the Switch's Upstream Port is handled. The Request can be blocked, redirected, or allowed to proceed directly to its target.

  If any Functions other than the Switch Upstream Port are associated with the Upstream Port, this field has no effect on accesses to their Memory BAR Space. Such access is controlled by the ACS Extended Capability (if present) in the Switch Upstream Port.

- ACS Unclaimed Request Redirect: must be implemented by Switch Downstream Ports that support ACS Enhanced Capability; is not applicable to Root Ports.

  When enabled, incoming Requests received by the Switch Downstream Port's Ingress and targeting Memory Space within the memory window of a Switch Upstream Port that is not within a memory window or Memory BAR Target of any Downstream Port within the Switch are redirected Upstream out of the Switch.

  When not enabled, such Requests are handled by the Switch as a UR.

- ACS I/O Request Blocking: must be implemented by Root Ports and Switch Downstream Ports that support ACS Enhanced Capability.

  When enabled, the Port must handle an Upstream I/O Request received by the Port's Ingress as an ACS Violation.

Add a Section 6.12.4:

6.12.4 ACS Enhanced Capability

ACS Enhanced Capability is an additional set of ACS control mechanisms to improve the level of isolation and protection provided by ACS. ACS Enhanced Capability defines the following additional access control mechanisms:

- ACS DSP Memory Target Access
- ACS USP Memory Target Access
- ACS Unclaimed Request Redirect

[2] This includes any Memory Space allocated by an Expansion ROM Base Address register (BAR). This also includes any Memory Space allocated by EA entries with a BEI value of 0, 1, 7, or 8. See Section 7.8.5.2.

- ACS I/O Request Blocking

Through these mechanisms, ACS Enhanced Capability provides protection and consistent handling of Requests directed toward regions not covered by the original ACS mechanisms.

Implementation Note: ACS Redirect and Guest Physical Addresses (GPAs)

ACS redirect mechanisms were originally architected to enable fine-grained access control for P2P Memory Requests, by redirecting selected Requests Upstream to the RC, where validation logic determines whether to allow or deny access. However, ACS redirect mechanisms can also ensure that Functions under the direct control of VMs have their DMA Requests routed correctly to the Translation Agent in the host, which then translates their guest physical addresses (GPAs) into host physical addresses (HPAs).

GPA ranges used for Memory Space vs. DMA are not guaranteed to coincide with HPA ranges, which the PCIe fabric uses for Memory Request routing and access control. If any GPAs used for DMA fall with within the HPA ranges used for Memory Space, legitimate or malicious packet misrouting can result.

ACS redirect mechanisms can ensure that Upstream Memory Requests with GPAs intended for DMA never get routed to HPA Memory ranges. ACS P2P Request Redirect handles this for (1) peer accesses between Functions within a Multi-Function Device and (2) peer accesses between Downstream Ports within a Switch or RC. ACS P2P Egress Control with redirect handles this in a more fine-grained manner for the same two cases.

Redirect mechanisms introduced with ACS Enhanced Capability handle this for additional cases. ACS DSP Memory Target Access with redirect handles this for Downstream Port Memory Resource ranges. ACS USP Memory Target Access with redirect handles this for Switch Upstream Port Memory Resource ranges. In Switches, ACS Unclaimed Request Redirect handles this for any areas within Upstream Port Memory apertures that are not handled by the other ACS redirect mechanisms.

Together these ACS redirect mechanisms can ensure that Upstream Memory Requests with GPAs intended for DMA are always routed or redirected to the Translation Agent in the host, and those with GPAs intended for P2P are still routed as originally architected.

…

Make the following change to Table 7-73:

!!!TODO: Update Figure 7-91

Table 7-73 ACS Capability Register

Bit Location	Register Description	Attributes
0	**ACS Source Validation ~~(V)~~ –** …	RO
1	**ACS Translation Blocking ~~(B)~~ –** …	RO

!!!TODO: Update Figure 7-91

Table 7-73 ACS Capability Register

Bit Location	Register Description	Attributes
2	**ACS P2P Request Redirect ~~(R)~~ –** …	RO
3	**ACS P2P Completion Redirect ~~(C)~~ –** …	RO
4	**ACS Upstream Forwarding ~~(U)~~ –** …	RO
5	**ACS P2P Egress Control ~~(E)~~ –** …	RO
6	**ACS Direct Translated P2P ~~(T)~~ –** …	RO
7	**ACS Enhanced Capability -** Required for Switch Downstream Ports that support the ACS Enhanced Capability mechanisms. If 1b, indicates that the component supports the following mechanisms: • ACS DSP Memory Target Access • ACS USP Memory Target Access • ACS Unclaimed Request Redirect • ACS I/O Request Blocking	RO
…	…	…

Make the following changes to Table 7-74:

Table 7-72 ACS Control Register

Bit Location	Register Description	Attributes
0	**ACS Source Validation Enable ~~(V)~~ –** …	RW
1	**ACS Translation Blocking Enable ~~(B)~~ –** …	RW
2	**ACS P2P Request Redirect Enable ~~(R)~~ –** …	RW
3	**ACS P2P Completion Redirect Enable ~~(C)~~ –** …	RW
4	**ACS Upstream Forwarding Enable ~~(U)~~ –** …	RW

Bit Location	Register Description	Attributes
5	**ACS P2P Egress Control Enable ~~(E)~~ –** …	RW
6	**ACS Direct Translated P2P Enable ~~(T)~~ –** …	RW
7	**ACS I/O Request Blocking bit –** if Set, Upstream I/O Requests received by the Downstream Port must be handled as ACS Violations. This field is required for Root Ports and Switch Downstream Ports if the ACS Enhanced Capability bit is Set. Must be RsvdP otherwise. The default value of this bit (if implemented) is 0b.	RW/RsvdP
9:8	**ACS DSP Memory Target Access–** This field determines how a Downstream Port handles Upstream Memory Requests attempting to access any Memory BAR Space on an applicable Root Port or Switch Downstream Port (including the Ingress Port). See Section 6.12.1.1. Defined Encodings are: 00b Direct Request access enabled 01b Request blocking enabled 10b Request redirect enabled 11b Reserved This field is required for Root Ports and Switch Downstream Ports if the ACS Enhanced Capability bit is Set and there is applicable Memory BAR Space to protect. Must be RsvdP otherwise. The default value of this field (if implemented) is 00b.	RW/RsvdP
11:10	**ACS USP Memory Target Access –** This field determines how a Switch Downstream Port handles Upstream Memory Requests attempting to access any Memory BAR Space on the Switch Upstream Port. See Section 6.12.1.1. Defined Encodings are: 00b Direct Request access enabled 01b Request blocking enabled 10b Request redirect enabled 11b Reserved This field is required for Switch Downstream Ports if the ACS Enhanced Capability bit is Set and there is applicable Memory BAR Space in the Switch Upstream Port to Protect. Must be RsvdP otherwise. The default value of this field (if implemented) is 00b.	RW/RsvdP

Bit Location	Register Description	Attributes
12	**ACS Unclaimed Request Redirect bit –** Determines how a Switch Downstream Port handles incoming Requests targeting Memory Space within the Memory aperture of the Switch Upstream Port that is not within a Memory aperture or Memory BAR Space of any Downstream Port within the Switch. When Set, the Switch must forward such Requests Upstream out of the Switch. When Clear, the switch must handle such Requests as an Unsupported Request (UR). This field is required for Switch Downstream Ports if the ACS Enhanced Capability bit is Set. Must be RsvdP otherwise. The default value of this bit (if implemented) is 0b.	RW/RsvdP